

# Requirements to Design to Code: Towards a Fully Formal Approach to Automatic Code Generation

Michael G. Hinchey<sup>1</sup>      James L. Rash<sup>1</sup>      Christopher A. Rouff<sup>2</sup>

November 30, 2004

<sup>1</sup> Information Systems Division Code 580 NASA Goddard Space Flight Center Greenbelt, MD 20771 {michael.g.hinchey, james.l.rash}@nasa.gov	<sup>2</sup> SAIC Suite 300 4001 N. Fairfax Drive Arlington, VA 22203 rouffc@saic.com
--	---

## Abstract

A general-purpose method to mechanically transform system requirements into a provably equivalent model has yet to appear. Such a method represents a necessary step toward high-dependability system engineering for numerous possible application domains, including sensor networks and autonomous systems. Currently available tools and methods that start with a formal model of a system and mechanically produce a provably equivalent implementation are valuable but not sufficient. The “gap” that current tools and methods leave unfilled is that their formal models cannot be proven to be equivalent to the system requirements as originated by the customer. For the classes of systems whose behavior can be described as a finite (but significant) set of scenarios, we offer a method for mechanically transforming requirements (expressed in restricted natural language, or in other appropriate graphical notations) into a provably equivalent formal model that can be used as the basis for code generation and other transformations.

**Key Words:** Validation, verification, formal methods, automatic code generation, sensor networks

## Contents

### 1 Introduction

2

<b>2</b>	<b>Problem Statement</b>	<b>2</b>
2.1	Specifications, Models, and Designs . . . . .	3
2.2	A Novel Approach . . . . .	5
<b>3</b>	<b>Technical Approach</b>	<b>6</b>
3.1	R2D2C . . . . .	6
3.2	Short-cut R2D2C . . . . .	9
<b>4</b>	<b>A Simple Example</b>	<b>10</b>
4.1	Specification of LOGOS . . . . .	11
<b>5</b>	<b>Application Areas</b>	<b>14</b>
<b>6</b>	<b>Related Work</b>	<b>15</b>
<b>7</b>	<b>Conclusions and Future Work</b>	<b>16</b>

# 1 Introduction

Sensor networks and other highly distributed autonomous systems cannot attain high dependability without addressing software dependability issues. Development of a system that will have a high level of reliability requires the developer to represent the system as a formal model that can be proven to be correct. Through the use of currently available tools, the model can then be automatically transformed into code with minimal or no human intervention to reduce the chance of inadvertent insertion of errors by developers. Automatically producing the formal model from customer requirements would further reduce the chance of insertion of errors by developers.

The need for ultra-high dependability systems increases continually, along with a correspondingly increasing need to ensure correctness in system development. By “correctness”, we mean that the implemented system is equivalent to the requirements, and that this equivalence can be proved mathematically.

Available system development tools and methods that are based on formal models provide neither automated generation of the models from requirements nor automated proof of correctness of the models. Therefore, today there is no automated means to produce a system or a procedure that is a provably correct implementation of the customer’s requirements. Further, requirements engineering as a discipline has yet to produce an automated, mathematics-based process for requirements validation.


## 2 Problem Statement

Automatic code generation from requirements has been the ultimate objective of software engineering almost since the advent of high-level programming languages, and calls for a “requirements-based programming” capability have become deafening [Har04]. Several tools and products exist in the marketplace for automatic code generation from a given model. However, they typically generate code, portions of which are never executed, or portions of which cannot be justified from either the requirements or the model. Moreover, existing tools do not and cannot overcome the fundamental inadequacy of all currently available automated development approaches, which is that they include no means to establish a provable equivalence between the requirements stated at the outset and either the model or the code they generate.


Traditional approaches to automatic code generation, including those embodied in commercial products such as Matlab [Mat00], in system development toolsets such as the B-Toolkit [LH96] or the VDM++ toolkit [IFA00], or in academic research projects, presuppose the existence of an explicit (formal) model of reality that can be used as the basis for subsequent code generation (see Figure 1 (a)). While such an approach is reasonable, the advantages and disadvantages of the various modeling approaches used in computing are well known and certain models can serve well to highlight certain issues while suppressing other less relevant details [Par95]. It is clear that the converse is also true. Certain models of reality, while successfully detailing many of the issues of interest to developers, can fail to capture some important issues, or perhaps even the most important issues. Existing reverse-engineering approaches suffer from a similar plight. In typical approaches, such as the one illustrated in Figure 1 (b), a model is extracted from an existing system and is then represented in various ways, for example as a digraph [McL92]. The re-engineering process then involves using the resulting representation as the basis for code generation, as above.

### 2.1 Specifications, Models, and Designs

The model on which automatic code generation is based is referred to as a design, or more correctly, a design specification. There is typically a mismatch between the design and the implementation (sometimes termed the “specification-implementation gap”) in that the process of going from a suitable design to an implementation involves many practical decisions that must be made by the automated tool used for code generation without any clear-cut justifications, other than the predetermined implementation decisions of the tool designers. There is a more problematic “gap”, termed the “analysis-specification gap”, that emphasizes the



**(a) traditional development process**



**(b) reverse engineering process**

**Figure 1: (a) traditional software development process from requirements to code, and (b) reverse engineering from code to a system description.**

problem of capturing requirements and adequately representing them in a specification that is clear, concise, and complete. This specification must be formal, or proof of correctness is impossible [Bau80]. Unfortunately, there is reluctance by many to embrace formal specification techniques, believing them to be difficult to use and apply [Hal90] [BH95], despite many industrial success stories [HB95] [HB99].

Our experience at NASA Goddard Space Flight Center (GSFC) has been that while engineers are happy to write descriptions as natural language scenarios, or even using semi-formal notations such as UML use cases, they are loath to undertake formal specification. Absent a formal specification of the system under consideration, there is no possibility of determining any level of confidence in the correctness of an implementation. More importantly, we must ensure that this formal specification fully, completely, and consistently captures the requirements set forth at the outset. Clearly, we cannot expect requirements to be perfect, complete, and consistent from the outset, which is why it is even more important to have a formal specification, which can highlight errors, omissions, and conflicts. The formal specification must also reflect changes and updates from system maintenance as well as changes and compromises in requirements, so that it remains an accurate representation of the system.

## 2.2 A Novel Approach

Our approach involves providing a mathematically tractable round-trip engineering approach to system development. The approach described herein is provisionally named R2D2C (“Requirements to Design to Code”).

In this approach, engineers (or others) may write specifications as scenarios in constrained (domain-specific) natural language, or in a range of other notations (including UML use cases). These will be used to derive a formal model (Figure ??) that is guaranteed to be equivalent to the requirements stated at the outset, and which will subsequently be used as a basis for code generation. The formal model can be expressed using a variety of formal methods. Currently we are using CSP, Hoare’s language of Communicating Sequential Processes [Hoa78] [Hoa85], which is suitable for various types of analysis and investigation, and as the basis for fully formal implementations as well as for use in automated test case generation, etc.

## 3 Technical Approach

R2D2C is unique in that it allows for full formal development from the outset, and maintains mathematical soundness through all phases of the development process, from requirements through to automatic code generation. The approach may also




Figure 2: **The R2D2C approach, generating a formal model from requirements and producing code from the formal model, with automatic reverse engineering.**

be used for reverse engineering, that is, in retrieving models and formal specifications from existing code, as shown in Figure ???. The approach can also be used to “paraphrase” (in natural language, etc.) formal descriptions of existing systems. In addition, the approach is not limited to generating high-level code. It may also be used to generate business processes and procedures, and we are currently experimenting with using it to generate instructions for robotic devices to be used on the Hubble Robotic Servicing Mission (HRSM). We are also experimenting with using it as a basis for an expert system verification tool, and as a means of capturing expert knowledge for expert systems. Such potential applications will be described in Section 5.

Section 3.1 describes the approach at a relatively high level. Section 3.2 describes an intermediate version of the approach for which we have built a prototype tool, and with which we have successfully undertaken some examples.

### 3.1 R2D2C

The R2D2C approach involves a number of phases, which are reflected in the system architecture described in Figure 3. The following describes each of these phases.

- D1** Scenarios Capture: Engineers, end users, and others write scenarios describing intended system operation. The input scenarios may be represented in a constrained natural language using a syntax-directed editor, or may be represented in other textual or graphical formss.
- D2** Traces Generation: Traces and sequences of atomic events are derived from the scenarios defined in D1.

- D3 Model Inference:** A formal model, or formal specification, expressed in CSP is inferred by an automatic theorem prover – in this case, ACL2 [KMM00] – using the traces derived in phase 2. A deep<sup>1</sup> embedding of the laws of concurrency [HJ95] in the theorem prover gives it sufficient knowledge of concurrency and of CSP to perform the inference. The embedding will be the topic of a future paper.
- D4 Analysis:** Based on the formal model, various analyses can be performed, using currently available commercial or public domain tools, and specialized tools that are planned for development. Because of the nature of CSP, the model may be analyzed at different levels of abstraction using a variety of possible implementation environments. This will be the subject of a future paper.
- D5 Code Generation:** The techniques of automatic code generation from a suitable model are reasonably well understood. The present modeling approach is suitable for the application of existing code generation techniques, whether using a tool specifically developed for the purpose, or existing tools such as FDR [FSE99], or converting to other notations suitable for code generation (e.g., converting CSP to B [But99]) and then using the code generating capabilities of the B Toolkit.

It should be re-emphasized that the “code” that is generated may be code in a high-level programming language, low-level instructions for (electro-) mechanical devices, natural-language business procedures and instructions, or the like. As Figure ?? illustrates, the above process may also be run in reverse:

- R1 *Model Extraction:*** Using various reverse engineering techniques [vZ93], a formal model expressed in CSP may be extracted.
- R2 *Traces Generation:*** The theorem prover may be used to automatically generate traces based on the laws of concurrency and the embedded knowledge of CSP.
- R3 *Analysis:*** Traces may be analyzed, used to check for various conditions, undesirable situations arising, etc.
- R4 *Paraphrasing:*** A description of the system (or system components) may be retrieved in the desired format (natural language scenarios, UML use cases, etc.).

---

<sup>1</sup>“Deep” in the sense that the embedding is semantic rather than merely syntactic.

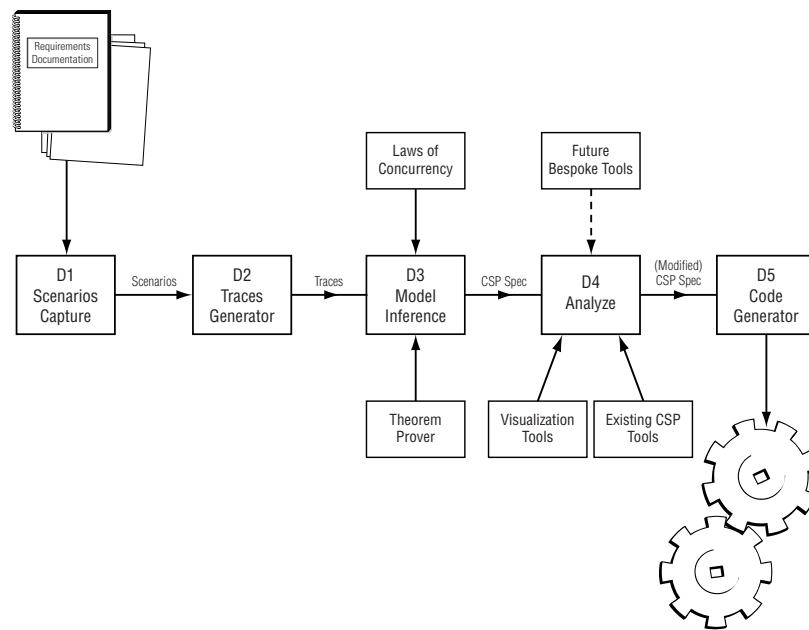


Figure 3: **The entire process with D1 thru D5 illustrating the development approach and R1 thru R4 the reverse engineering.**



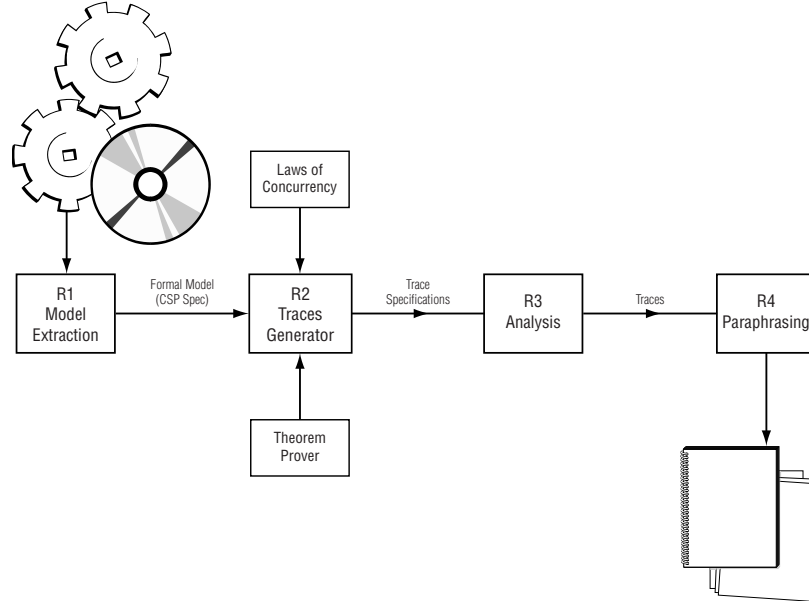


Figure 4: **Reverse engineering of system using R2D2C.**

Paraphrasing, whereby more understandable descriptions (above and beyond existing documentation) of existing systems or system components are extracted, is likely to have useful application in future system maintenance for systems that have lost the original design documents or that have been modified so much that the original design and requirements document do not reflect the current system.

### 3.2 Short-cut R2D2C

The approach described in Section 3.1 is the way that R2D2C is intended to be applied, from requirements specification through to code generation. However, the approach requires significant computing power in the form of an automated theorem prover performing significant inferences based on traces input and its “knowledge” of the laws of concurrency. While this is well warranted for certain applications, it is likely to be beyond the resources of many developers and organizations. As a practical concession, we also define a reduced version of R2D2C called the shortcut version (Figure 5), whereby the use of a theorem prover is avoided, yet without sacrificing high confidence in the validity of the approach. The following describes each of the phases for the shortcut R2D2C:

- S1** *Scenarios Capture*: As before, intended system behavior is described by scenarios input in natural language, or an appropriate graphical or semi-formal notation.
- S2** *Translation to Intermediate Notation*: Scenarios are translated to an intermediate notation, termed EzyCSP, which is a simple natural language-like subset of CSP that can be used to describe a large number of situations and scenarios (recall that scenarios are domain specific).
- S3** *Analysis*: While far more simple than CSP, EzyCSP allows some simple analyses to be performed.
- S4** *Implementation in Java*: EzyCSP is sufficiently simple that it may easily be translated to Java and executed.

This simplified or short-cut approach clearly has significant disadvantages when compared to our full approach. Firstly, the correctness of the development process is contingent on the correctness of both the translation of scenarios to the intermediate (EzyCSP) notation and the translation of EzyCSP to Java. However, the correctness of the translators for these is assured via a proof of correctness undertaken with the ACL2 theorem prover. Secondly, we do not have a reverse process, suitable to support reverse and (ultimately) re-engineering, for free. However, a Java-to-EzyCSP translator would certainly be possible for highly constrained subsets of Java.

The significant advantage of this simplified approach, however, is that although a proof of correctness involving a theorem prover is still required, this is required exactly once and would be performed by the support system developers (presumably expert in the art). This is significantly less expensive computationally than using a theorem prover in the development of each individual application.

## 4 A Simple Example

The Lights-Out Ground Operating System (LOGOS) is a proof-of-concept NASA system for automatic control of ground stations when satellites pass overhead and under their control. The system exhibits both autonomous and autonomic properties [TRRH04] [THRR05], and operates by having a community of distributed autonomous software modules work cooperatively to perform the functions previously undertaken by human operators using traditional software tools, such as orbit generators and command sequence planners. A post-implementation formal specification of the system was undertaken in CSP. The interested reader is directed

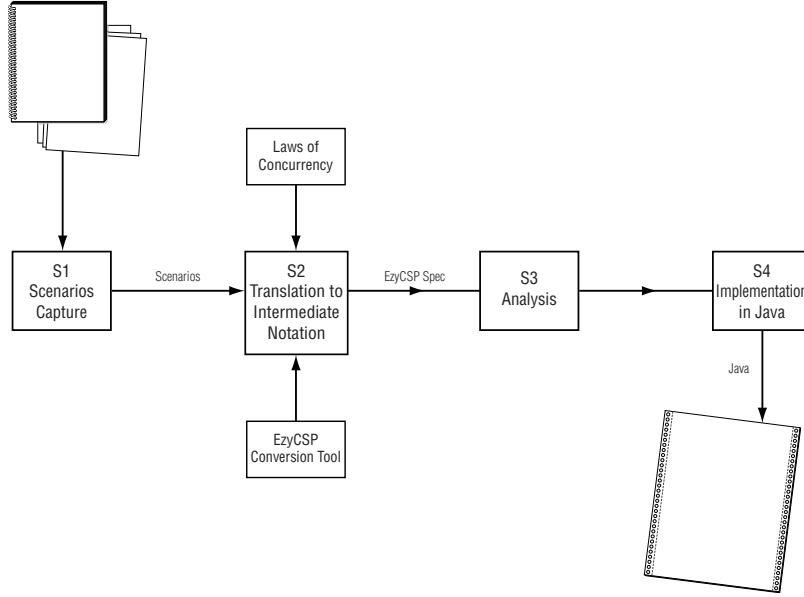


Figure 5: **Short cut R2D2C.**

to [?] [RRH00] and [HRR01] for a detailed discussion of our experiences in applying formal methods to LOGOS. Using CSP, a number of anomalies, conflicts, and omissions in the system were discovered that had not been detected in testing and/or actual execution. This experience is typical of highly distributed systems, such as sensor networks or other multi-agent based systems where dependability is both very important and very difficult to evaluate. The same approach can be used for space based WSN systems where a control station is in charge of several WSNs located on spacecrafts in deep space. An example is the Autonomous Nano Technology Swarm mission (ANTS) [CMN<sup>+</sup>00], which is at the concept development phase. This mission will send 1,000 pico-class (approximately 1 kg) spacecraft to explore the asteroid belt. The ANTS spacecraft will act as a sensor network making observations of asteroids and analyzing their composition.

#### 4.1 Specification of LOGOS

We will not consider the entire LOGOS/ANTS related system here. Although a relatively small system, it is too extensive to illustrate in its entirety in this paper. Instead, we will take a couple of example agents from the system, and illustrate their mapping from natural language descriptions through to simple Java imple-

mentations.

Let us first illustrate, via a trivial example, how scenarios map to CSP. Suppose we have the following as part of one of the scenarios for the system:

if the WSN Monitoring Agent receives a *fault* advisory from the  
WSN the agent sends the fault to the Fault Resolution Agent  
OR  
if the WSN Monitoring Agent receives engineering data from the WSN  
the agent sends the data to the Trending Agent

That part of the scenario could be mapped to structured text as:

*inWSNMA?fault* from WSN  
then *outWSNMA!fault* to FIRE  
else  
*inengWSNMA?data* from WSN  
then *outengWSNMA!data* to TREND

The laws of concurrency would allow us to derive the traces as:

$$\begin{aligned} tWSNMA \supseteq & \{ \langle \rangle, \langle inWSNMA, fault \rangle, \langle inWSNMA, fault, outWSNMA, fault \rangle \} \\ \cup & \{ \langle \rangle, \langle inengWSNMA, data \rangle, \langle inengWSNMA, data, outWSNMA, data \rangle \} \end{aligned}$$

From which, we can infer an equivalent CSP process specification as:

$$\begin{aligned} WSNMA = & inWSNMA?fault \rightarrow (outWSNMA!fault \rightarrow SKIP) \\ & | (inengWSNMA?data \rightarrow outengWSNMA!data \rightarrow SKIP) \end{aligned}$$

Let us now consider a slightly larger example, the Pager Agent, and illustrate its implementation in Java. The pager agent sends pages to engineers and controllers when there is a WSN anomaly and there is no analyst logged on to the system. The pager agent receives requests from the user interface agent that no analyst is logged on, gets paging information from the database agent that keeps relevant information about each user of the system (in this case the analyst's pager number), and, when instructed by the user interface agent that the analyst has logged on, stops paging. These scenarios can be restated in more structured natural language as follows:

if the Pager agent receives a request from the User Interface agent, the  
Pager agent sends a request to the database agent for an analyst's  
pager information and puts the message in a list of requests to the  
database agent  
OR

if the Pager agent receives a pager number from the database agent,  
then the pager agent removes the message from the paging queue  
and sends a message to the analyst's pager and adds the analyst  
to the list of paged people

OR

if the Pager agent receives a message from the user interface agent  
to stop paging a particular analyst, the pager sends a stop-paging  
command to the analyst's pager and removes the analyst from the  
paged list

OR

if the Pager agent receives another kind of message, reply to the sender  
that the message was not recognized

The above scenarios would then be translated into CSP. The following is a  
partial CSP description of the pager agent:

```
PAGER_BUSdb_waiting, paged = pager.In?msg →  
  case  
    GET_USER_INFOdb_waiting, paged, pagee, text  
      if msg = (START_PAGING, specialist, text)  
  
    BEGIN_PAGINGdb_waiting, paged, in_reply_to_id(msg), pager_num  
      if msg = (RETURN_DATA.pager_num)  
  
    STOP_CONTACTdb_waiting, paged, pagee  
      if msg = (STOP_PAGING, pagee)  
  
    pager.Iout(head(msg), UNRECOGNIZED)  
      → PAGER_BUSdb_waiting, paged  
    otherwise
```

This specification states that the process *PAGER\_BUS* receives a message on its “In” channel and stores it in a variable called “msg”. Depending on the contents of the message, one of four different processes is executed. If the message has a *START\_PAGING* performative, then the *GET\_USER\_INFO* process is called with parameters of the type of specialist to page (pagee) and the text to send the pagee. If the message has a *RETURN\_DATA* performative with a pagee's pager number, then the database has returned a pager number and the *BEGIN\_PAGING* process is executed with a parameter containing the original message id (used as a key to the db\_waiting set) and the passed pager number. The third type of message

that the pager agent might receive is one with a `STOP_PAGING` performative. This message contains a request to stop paging a particular specialist (stored in the `pagee` parameter). When this message is received, the `STOP_PAGING` process is executed with the parameter of the specialist type. If the pager agent receives any other message than the above three messages, an error message is returned to the sender of the message (which is the first item of the list) stating that the message is “UNRECOGNIZED”. After this, the `PAGER_BUS` process is again executed.

## 5 Application Areas

The motivation for this work was the need for automatic code generation for ultra high dependability systems, but the method described in this paper applicable in a number of other areas. The following areas may find significant value from the use of this system.

### *Sensor Networks —*

NASA is currently conducting research and development on sensor networks for planetary and solar system exploration as well as to support its Mission to Planet Earth. In addition to the ANTS mission, a similar mission is being considered to explore the rings of Saturn. Sensor networks are also being considered for planetary (e.g., Martian) exploration, to yield valuable scientific information on weather and geological aspects. For the Mission to Planet Earth, sensor networks are already being researched and developed towards capabilities for early warnings about natural disasters and climate change. With the nature of systems of sensor networks, the inter-relatedness of these systems all networked together will create a level of complexity that will require a new level of dependability and a corresponding new approach to system and software development.

Projected NASA sensor networks are highly distributed autonomous systems that must operate with a high degree of reliability, especially the solar system and planetary exploration networks, which will be out of touch with the Earth and mission control for long periods of time and must operate under extremes of dynamic environmental conditions. Due to the complexity of these systems as well as their distributed and parallel nature, they will have an extremely large state space and will be impossible to test using traditional testing techniques. The more “code” or instructions that can be generated automatically from a verifiably correct model, the less likely that people will introduce errors during the development process. In addition, the higher the level of abstraction that developers can work from, as is afforded through the use of scenarios to describe system behavior, the less likely that a mismatch will occur between requirements and implementation and the more likely that the system will be validated. Working from a higher level of abstraction

will also allow errors in the system to be more easily caught, since developers can better see the “big picture” of the system. In addition to allowing complex systems developers to work at a higher level of abstraction, R2D2C also converts the scenarios into a formal model that can be analyzed for concurrency-related errors as well as consistency and completeness, in addition to domain-specific errors.

*Expert Systems —*

We have been studying the potential use of this approach in the development, maintenance, and verification of expert systems. In particular, we have been giving consideration to using the R2D2C method in verifying the expert system used in the NASA ground control center for the POLAR spacecraft, which performs multi-wavelength imaging of the Earth’s aurora. The POLAR ground control expert system has rules written in the production system CLIPS [Gia92] for automated “lights out” (untended) operation of the spacecraft. A suitable translator from CLIPS, rather than natural language, to CSP (or EzyCSP) enables us to use this technology to examine existing expert system rule bases for consistency, etc. What has proven to be of great interest, however, is the ability to generate CLIPS rules from CSP (or EzyCSP), just as we would generate code in Java or C++. POLAR ground control center personnel expect this would be a great benefit because it would give them a means of capturing expert knowledge, from natural language description through to CLIPS rules, while maintaining correctness, which heretofore has not been available.

*Robotic Operations —*

As pointed out earlier, the “code” generated by this approach need not be specifically code in a programming language. To this end, we have been experimenting with generating code to control robotic devices. Perhaps more interesting is the use of this approach to investigate the validity and correctness of procedures for complex robotic assembly or repair tasks. We have begun exploratory work in this direction, to validate procedures from the Hubble Robotic Servicing Mission (HRSM) – for example, the procedures for replacement of cameras, etc. on the Hubble Space Telescope (HST).

## **6 Related Work**

Harel [Har01] [HM03] has advocated scenario-based programming through UML use cases and play-in scenarios. This work differs in that it uses scenarios in the form of structured text that is easily understandable by engineers and non-engineers. In addition, the results of converting the structured text to traces and then from traces to a formal model allows us to use a wide range of formal methods tools (e.g., model checkers), which can be used to verify and validate the system.

NASA Ames has been working on the automatic translation of UML use cases to executable code, and report success in using the approach on large applications [WSK03]. Our approach is different, however, in that we are not limited to UML use cases, nor to natural language. R2D2C will work equally well with any input mechanism whereby requirements can be represented as scenarios, and traces extracted. Our approach works equally well with graphical, mathematical, and textual requirements representations. More importantly, the key to our approach and what makes it invaluable for high-dependability applications is the full formal basis, and complete mathematical tractability from requirements through to code. To our knowledge, no other currently available automated development methodology can make this claim.

## 7 Conclusions and Future Work

R2D2C is a unique approach to the automatic derivation of ultra-high dependability systems. It is unique in that it supports fully (mathematically) tractable development from requirements elicitation through to automatic code generation (and back again). While other approaches have supported various subsets of the development lifecycle, there has been heretofore a “jump” in deriving from the requirements the formal model that is a prerequisite for sound automatic code generation. Yet, R2D2C is a simple approach, combining techniques and notations that are well understood, well tried and tested, and trusted. The novelty of the approach, and the part of the approach that achieves continuity in the development process, is the use of a theorem prover to reverse the laws of concurrency, and to achieve levels of inference that would be impossible for a human being to perform on all but trivial systems.

It is our contention that R2D2C, and other approaches that similarly provide mathematical soundness throughout the development lifecycle will:

- Dramatically increase assurance of system success by ensuring
  - completeness and consistency of requirements
  - that implementations are true to the requirements
  - that automatically coded systems are bug-free; and that
  - that implementation behavior is as expected
- Decrease costs and schedule impacts of ultra-high dependability systems through automated development
- Decrease re-engineering costs and delays



Future work will include improving the quality of the embedding of CSP in ACL2, and optimizing that for efficiency. We plan a plethora of support tools to allow us to easily change the level of abstraction in a formal model, to visualize various system models and changes in those models, and to aid in tracking changes through the development process (or the reverse engineering process). We plan to enhance our existing prototype to support the full version of R2D2C, to make it into a fully functional robust prototype, and to apply it to more significant examples than the one presented in this paper.

### Acknowledgements

This work was funded in part by the NASA Goddard Space Flight Center Technology Transfer Office. We are grateful to Ted Mecum, Nona Cheeks, Chris Kirkman, Diana Cox, Yvette Conwell-Brown, and Keith Dixon for their support and encouragement. Denis Gracanin (Virginia Tech) and John Erickson (University of Texas at Austin) worked with us on the intermediate approach described in Section 3.2, and undertook the implementation of the prototype described in that section. The approach described in this paper is protected under United States and international Patent Applications assigned to the United States government.

### References

- [Bau80] F. L. Bauer. A trend for the next ten years of software engineering. In H. Freeman and P. M. Lewis, editors, *Software Engineering*, pages 1–23. Academic Press, 1980.
- [BH95] J. P. Bowen and M. G. Hinchey. Seven more myths of formal methods, 1995.
- [But99] M. Butler. csp2b: A practical approach to combining csp and b. Technical Report DSSE-TR-99-2, Department of Electronics and Computer Science, University of Southampton, 1999.
- [CMN<sup>+</sup>00] S. A. Curtis, J. Mica, J. Nuth, G. Marr, M. Rilee, and M. Bhat. Ants (autonomous nano-technology swarm): An artificial intelligence approach to asteroid belt resource exploration. In *Proc. Int’l Astronautical Federation, 51st Congress*, October 2000.
- [FSE99] Ltd. Formal Systems (Europe). *Failures-Divergences Refinement: User Manual and Tutorial*. Formal Systems (Europe), Ltd., 1999.
- [Gia92] J. C. Giarratano. *CLIPS User’s Guide*. NASA Johnson Space Center, Houston, Texas, USA, 1992.

- [Hal90] Hall. Seven myths of formal methods, 1990.
- [Har01] D. Harel. From play-in scenarios to code: An achievable dream, 2001.
- [Har04] D. Harel. Requirements-based programming. In *Proc. ISOLA-2004*, 2004.
- [HB95] M. G. Hinchey and J. P. Bowen, editors. *Applications of Formal Methods*. Series in Computer Science. Prentice Hall International, Englewood Cliffs and Hemel Hempstead, 1995.
- [HB99] M. G. Hinchey and J. P. Bowen, editors. *Industrial-Strength Formal Methods in Practice*. FACIT Series. Springer-Verlag, London and New York, 1999.
- [HJ95] M. G. Hinchey and S. A. Jarvis. *Concurrent Systems: Formal Development in CSP*. International Series in Software Engineering. McGraw-Hill International, UK, London and New York, 1995.
- [HM03] D. Harel and R. Marelly. *Come, Let's Play: Scenario-Based Programming Using LSCs and the Play-Engine*. Springer-Verlag, 2003.
- [Hoa78] C. A. R. Hoare. Communicating sequential processes, 1978.
- [Hoa85] C. A. R. Hoare. *Communicating Sequential Processes*. International Series in Computer Science. Prentice Hall, Englewood Cliffs and Hemel Hempstead, 1985.
- [HRR01] M. G. Hinchey, J. L. Rash, and C. A. Rouff. Verification and validation of autonomous systems. In *Proc. SEW-26, 26th Annual NASA/IEEE Software Engineering Workshop*, Greenbelt, MD, November 2001. IEEE Computer Society Press.
- [IFA00] IFAD. *The VDM++ Toolbox User Manual*, 2000.
- [KMM00] M. Kaufmann, P. Manolios, and J S. Moore. *Computer-Aided Reasoning: An Approach*. Advances in Formal Methods Series. Kluwer, Boston, 2000.
- [LH96] K. Lano and H. Haughton. *Specification in B: an Introduction using the B-Toolkit*. Imperial College Press, 1996.
- [Mat00] Mathworks, Inc. *Getting Started with MATLAB*, 2000.

- [McL92] F. McLoughlin. Adaggio - an automated directed graph diagramming tool, 1992.
- [Par95] D. L. Parnas. *Applications of Formal Methods*, chapter Using Mathematical Models in the Inspection of Critical Software, pages 17–31. International Series in Computer Science. Prentice Hall, 1995.
- [RRH00] C. A. Rouff, J. L. Rash, and M. G. Hinchey. Experience using formal methods for specifying an agent-based system. In *Proc. ICECCS 2000, 6th IEEE International Conference on Engineering of Complex Computer Systems*, Tokyo, Japan, 2000. IEEE Computer Society Press.
- [THRR05] W. F. Truszkowski, M. G. Hinchey, J. L. Rash, and C. A. Rouff. Autonomous and autonomic systems: A paradigm for future space exploration missions, 2005.
- [TRRH04] W. Truszkowski, J. Rash, C. Rouff, and M. Hinchey. Some autonomic properties of two legacy multi-agent systems – logos and act. In *Proc. 11th IEEE International Conference on Engineering Computer-Based Systems (ECBS), Workshop on Engineering Autonomic Systems (EASe)*, pages 490–498, Brno, Czech Republic, May 24-27 2004. IEEE Computer Society Press.
- [vZ93] H. J. van Zuylen. *The REDO Compendium: Reverse Engineering for Software Maintenance*. John Wiley and Sons, 1993.
- [WSK03] J. Whittle, J. Saboo, and R. Kwan. From scenarios to code: An air traffic control case study. In *Proc. ICSE-25, 25th IEEE/ACM International Conference on Software Engineering*, pages 490–495, Portland, OR, USA, 2003. IEEE Computer Society Press.